# Run IIb CDF Detector Project

# Completion Report

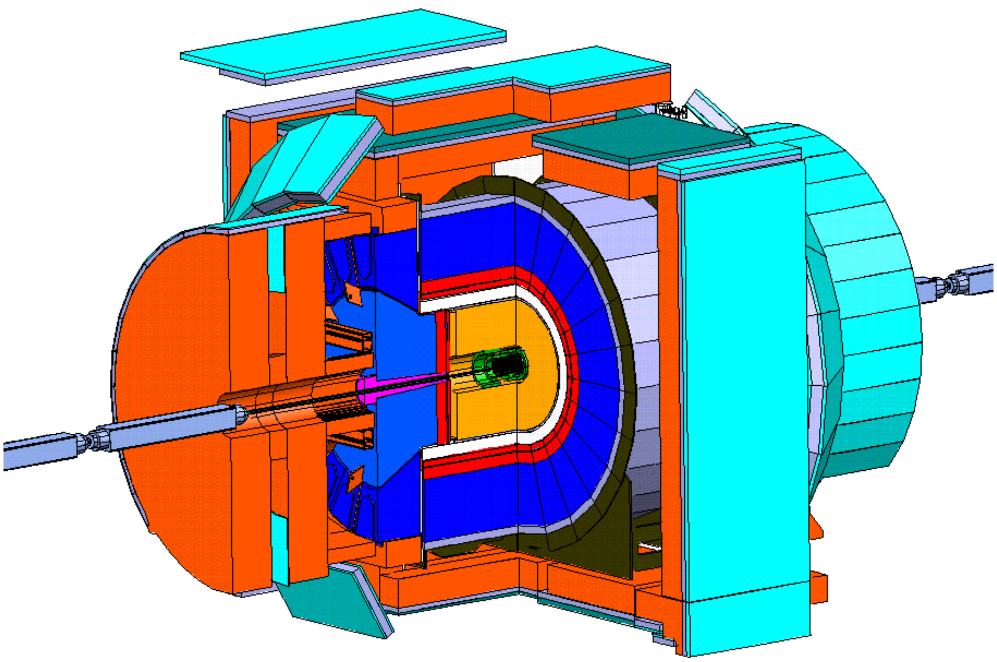

**July 2006**

# Table of Contents

## Introduction:

This report documents the completion of the Run IIb CDF Detector Project at Fermi National Accelerator Laboratory (Fermilab). This project was completed under budget and ahead of schedule. The final DOE MIE cost for the project was $7.2 million. The project was able to achieve its mission need and also return $1.0 million of unused contingency, in addition to $2.2 million that was returned in June 2005 to the Office of High Energy Physics for other uses. Essentially all costs were paid by May 2006, well ahead of the CD-4 date of November 2006. Information on the technical deliverables, cost performance, and schedule performance is included in this report.

## Historical Background

CDF first detected $\bar{p}$ p collisions in 1985. The detector collected data in 1987, 1988-89, 1992-93 ("Run Ia") and 1994-1996 ("Run Ib"). Collider Run IIa began in 2001. A large number of important physics results have been produced by CDF and have been published in numerous articles in refereed physics journals. These results include the discovery of the top quark and precision measurements of its mass and production cross section, precision measurement of the W boson mass; a broad program of electroweak measurements; QCD measurements; B physics, including measurement of lifetimes of exclusive states; and Exotic Physics including limits on the production of a variety of non-Standard Model objects.

CDF has gone through periods of extensive upgrades. Between 1989 and 1992, the detector was improved in several ways. This included the addition of a silicon vertex detector; additional muon detectors to increase the muon acceptance; improvements to existing muon systems, and a new inner tracking chamber used to measure the *z* position of event vertices. The experiment recorded 110 $pb^{-1}$ of integrated luminosity during the 1992-96 operating period (Run I).

In October 1990 a proposal was submitted to upgrade the CDF detector to allow it to continue to exploit the physics opportunities as improvements, including the Main Injector, were made to the Fermilab collider. The running conditions for collider Run II specified that the detector must be capable of handling peak luminosity up to $2 \times 10^{32}$ $cm^{-2}$ $sec^{-1}$, bunch spacing as small as 132 ns, and an integrated luminosity of 2 $fb^{-1}$. The CDF Run IIa upgrade included replacing the plug and forward gas calorimeters with a new scintillator-based calorimeter and replacing the Central Tracking Chamber with a device with shorter drift time to allow tracking in a high-luminosity environment. A completely new silicon system was built and installed. The front-end electronics and trigger systems were upgraded to accommodate data-taking at higher rates and with shorter bunch spacing. Muon detection systems were upgraded to increase acceptance and allow the electronics to work with shorter bunch spacing. The data acquisition system was upgraded to increase throughput and reliability. A new time-of-flight detector was added, as were new detectors in the forward region. The CDF Upgrade Project for Run IIa was successfully completed in March 2001.

The Run IIb CDF Detector Project was submitted in December 2002 for the purpose of extending the useful life of the CDF detector to instantaneous luminosities of 3 x $10^{32}$ $cm^{-2}$ $sec^{-1}$(with 396 ns crossing separation) and an integrated luminosity of 15 $fb^{-1}$. The integrated luminosity requirement was revised to 8 $fb^{-1}$ during the construction of the project, and the technical scope was adjusted accordingly. The project that remained after this change included upgrades to the data acquisition to accommodate the high event rate, upgrades to the track triggers to address the complexity of the events seen at high luminosity, and upgrades to the calorimeter system that replaced gas chamber systems with scintillator.

**Scope**

The Run IIb CDF Detector Project prolongs the useful life of the Collider Detector at Fermilab (CDF) for operation at higher luminosity than expected in the original design. Specifically, the detector must be capable of handling peak luminosity up to 3 x $10^{32}$ $cm^{-2}$ $sec^{-1}$. Several detector systems required replacement or modification in order to meet these requirements.

The major tasks of this upgrade were:

- Replace the silicon micro-vertex detector with a device capable of withstanding the expected radiation dose for Run IIb and with fast $r$-$\phi$ (axial) and small angle stereo readout. This portion of the project was cancelled due to decreased accelerator performance expectations.

- Upgrade the calorimeter by replacing the Central Preradiator Chamber with a device with shorter response time to allow operation in a high-luminosity environment, and adding timing information to the electromagnetic calorimeters.

- Upgrade the data acquisition and trigger systems to increase throughput needed for higher luminosity operation and efficiently trigger on the higher multiplicity events of Run IIb.

CD-4 definitions contained in the Run IIb CDF Detector Project Execution Plan were met and are included in the detailed descriptions of the subprojects presented below organized by Work Breakdown Structure element. The installation activities for the Run IIb Detector were not part of this project but were managed in a similar fashion using project management tools. All components of the project have either been included into operations, or are being actively commissioned with beam.

## 1.1 Silicon Detector Upgrade

Collaborating Institutions:

University of Tsukuba, Japan
Universita' di Padova, Italy
Johns Hopkins University

Duke University
Purdue University
University of New Mexico
Fermilab
University of California, Davis
University of Illinois at Urbana-Champaign
Lawrence Berkeley Laboratory
Carnegie Mellon University
University of California, Santa Barbara
Wayne State University
Pohjois-Savo Polytechnic, Finland
Okayama University, Japan
Center for High Energy Physics, Kyungpook National University, South Korea
University of Helsinki, Finland
Acedemia Sinica, Taiwan
Osaka City University, Japan
SungKyunKwan University, South Korea
Seoul National University, South Korea
Universita' di Bologna, Italy

The silicon detector upgrade was motivated by the integrated radiation dose that was predicted for the detector currently in use and the effect this dose would have on performance. Significant reductions in efficiency were expected after the delivered luminosity exceeded 4 $fb^{-1}$. The upgraded replacement detector was developed to the point where several of its components were ready for production readiness review, and final production assemblies were being developed. The integrated luminosity projections were reduced in 2002 and 2003, which motivated the Fermilab Director to recommend cancellation of full production of upgraded silicon detectors for both CDF and DØ. Consequently, the project scope was reduced in December 2003 to eliminate the full production of the detector.

The scope that remained included the construction of a demonstration device that contained approximately seven percent of the detector elements of the full device. This demonstration device was constructed with production parts and tested for electronics readout, noise levels, charge collection, etc. Several of the components developed for the silicon detector are illustrated below.

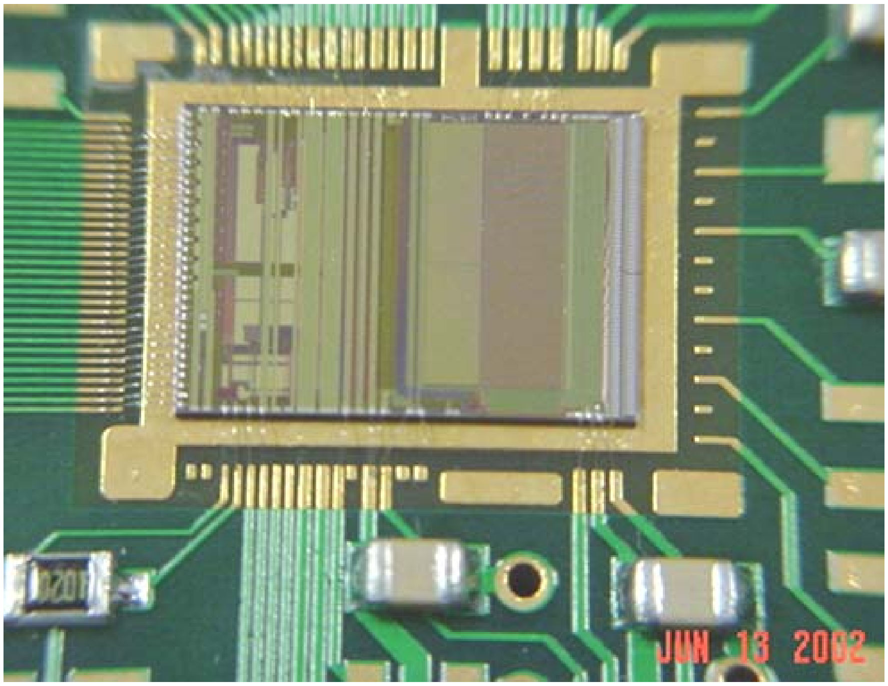


**SVX4 Chip**

The front end readout chip developed for the project was named SVX4 and was manufactured in 0.25 μm CMOS technology. The small feature size of these circuits gave them an inherent radiation hardness that surpassed the SVX3 chips that readout the existing silicon detector. This single production run of this chip was made, which provided good yield for the chips that were fully processed. This device has since been shared with groups at both Jefferson National Laboratory and Brookhaven National Laboratory for readout of silicon detectors at those facilities.

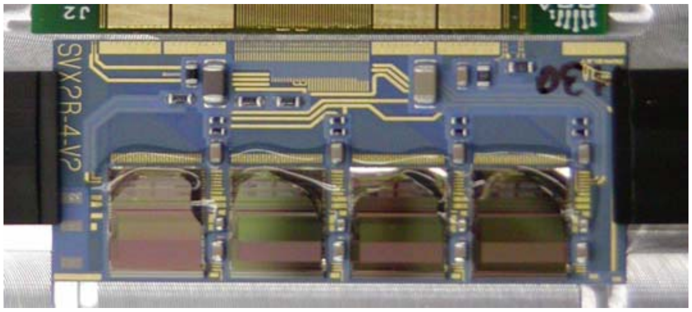

**SVX IIb Hybrid**

The SVXIIb Hybrid is a small circuit board that is fastened directly to the silicon sensors. Each hybrid contains four SVX4 chips and reads out a pair of silicon sensors. Low mass cables took the signals from the hybrids to the next connection point in the data acquisition, the Junction Port Cards.

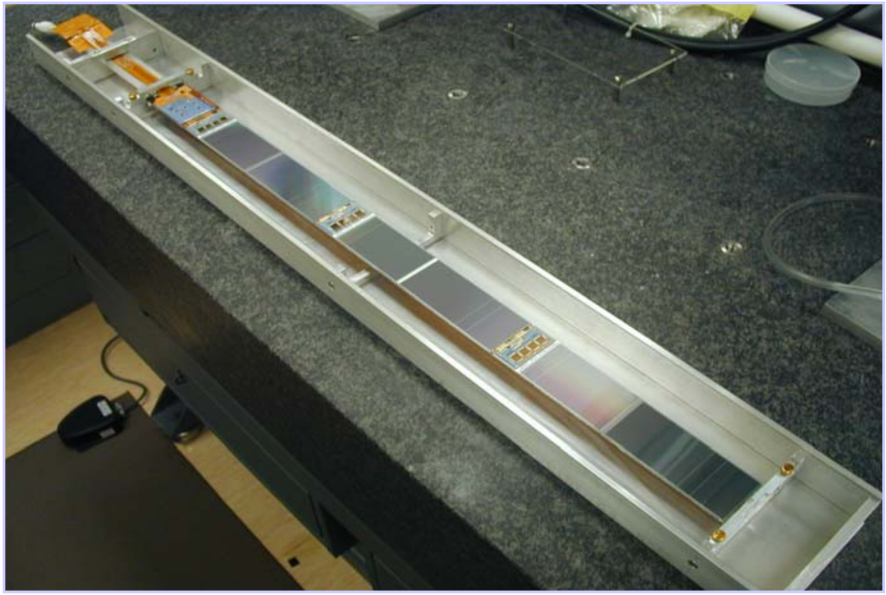

**SVX IIb Stave**

The basic building block of the SVX IIb silicon detector was a unit known as the stave. Each stave consisted of six modules of one hybrid and two sensors each.  The stave held three such modules on each side.  One side held sensors that were oriented to provide axial information on track positions, and the opposite side provided a shallow angle stereo measurement.  The full detector was to consist of 180 identical staves, which were used for all but the innermost layer.  The layer of silicon closest to the beam required a smaller device than could be economically constructed from a stave A special carbon fiber structure was built to hold axial measuring sensors and their hybrids for this region.

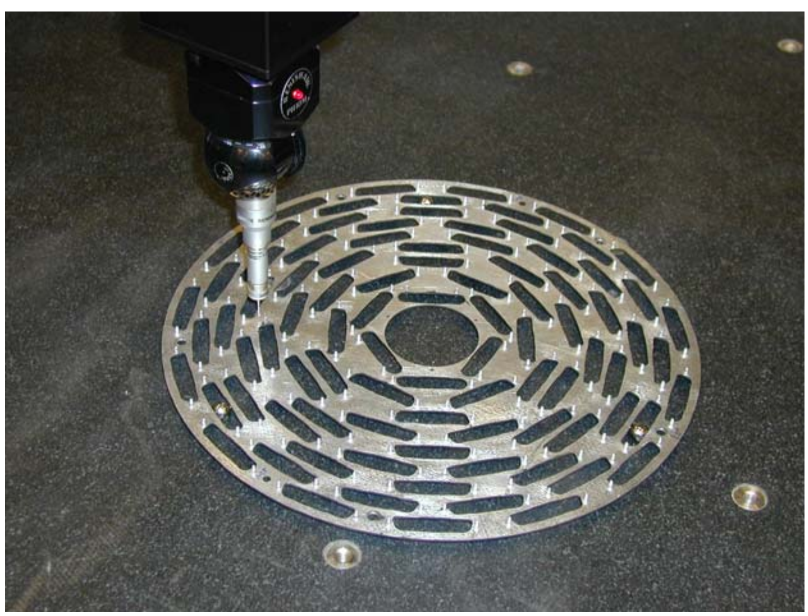

**SVX IIb Outer Bulkhead**

The full detector was designed to cover the tracking volume with two identical structures, referred to as "barrels." Each barrel was built from 90 staves, whose positions were held by two lightweight bulkheads.

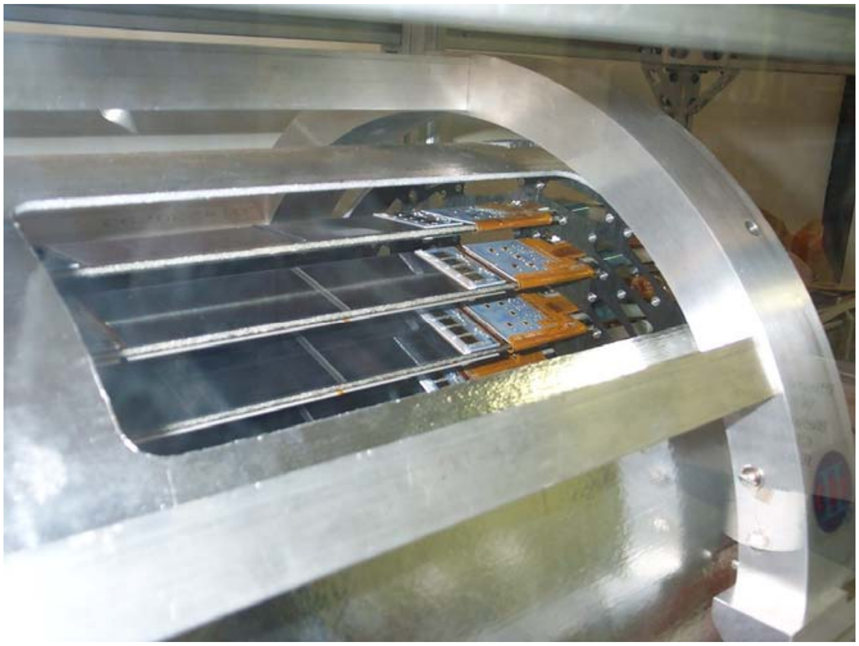

**The SVXIIb Detector**

The final demonstration detector is shown in the photo above. This assembly contains 15 staves, covering approximately 60° in the transverse view. The stave mounts to the bulkhead can be seen, as well as the hybrid and sensor pair of the endmost modules for the visible staves. A detailed description of performance has recently been accepted for published.[1]

### 1.2.1 Central Preradiator Upgrade

Collaborating Institutions:

- Argonne National Laboratory
- Fermilab
- INFN (Pisa, Rome, Trieste/Udine)
- JINR (Dubna)
- Kyungpook National University
- Michigan State Unversity
- Rockefeller University
- University of Tsukuba

The CDF Central Preshower (CPR) detector improves electron identification and provides the only photon background subtraction for photons above 35 GeV. The Central Crack detector (CCR) tags high energy photons that would otherwise be lost in the crack between calorimeter wedges. The upgrades of these detectors were motivated by the expected large occupancies of high luminosity running. The previous gas detectors integrated charge over four beam crossings, and were replaced by faster scintillator detectors that collected charge within one beam crossing. In addition, the new detectors have much better segmentation and fiducial coverage.

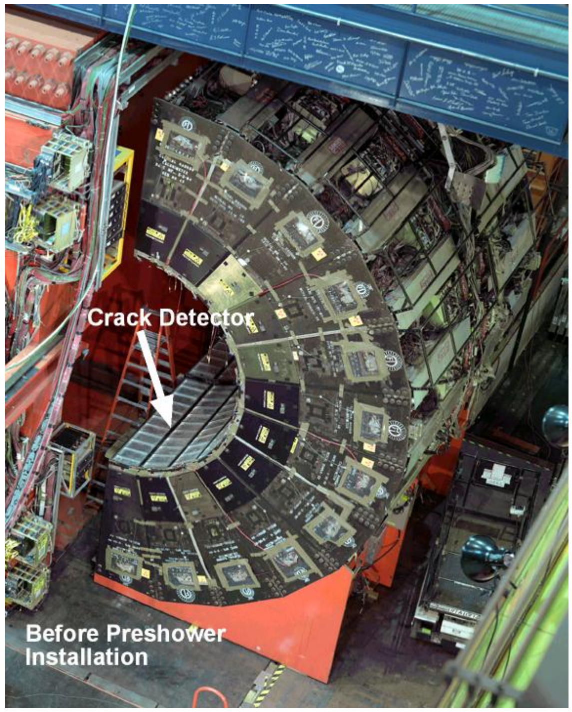


**CDF Detector Showing Crack And Preshower Locations**

An extensive R&D program into scintillator, fibers, fiber preparation, scintillator-groove design, fiber-optical cables, photomultipliers (PMTs), and mechanical design was carried out by Argonne National Laboratory (ANL), Fermilab , INFN, JINR (Dubna), Michigan State University, Rockefeller University, and Tsukuba University. Since the CPR detector needed to be fully efficient for minimum-ionizing particles (MIPs), the thickest possible scintillator that would fit in the available space was selected, 2cm Dubna scintillator. The CCR detector used 5mm thick Bicron BC408 scintillator. Kuraray fibers were used for both CPR and CCR modules, while the optical fibers leading to the PMTs were made by PolHiTech. The wavelength-shifting fibers were cut, polished, mirrored, and spliced to clear fibers. The scintillator groove design is a double-loop keyhole groove developed at Rockefeller. The optical path consisted of two parts, the spliced fibers leading from the scintillator to a connector on the end of the module, and clear fiber optical cables leading to the PMTs. The photomultipliers used are 16-channel multi-anode Hamamatsu R5900 PMTs.

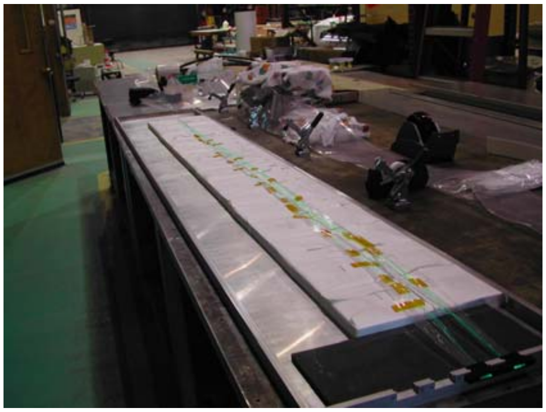 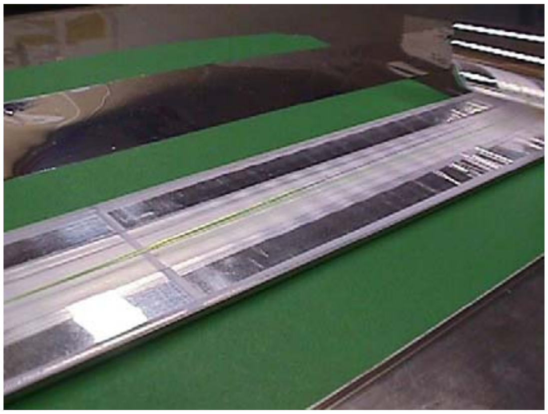

**Preshower Detector Construction** **Crack Detector Construction**

Assembly of the CPR proceeded as follows: the scintillator tiles were cut and polished in Dubna, then shipped to Fermilab for groove cutting, then to ANL for final assembly. The fibers were prepared at Fermilab, shipped to MSU for optical cable assembly, then to ANL. Aluminum outer shells were cut and assembled at ANL. After all the parts were compiled, teams of technicians from INFN, Korea, and Rockefeller, about 40 technicians total spread over six months, assisted ANL personnel in the final detector assembly and testing. The CCR assembly proceeded in a similar way at the same time.

**Calorimeter Installation**

Installation was challenging, but went very smoothly. Pieces of a three-story scaffold had to be passed thru a 13” x 17” hole and assembled inside the detector. This had to happen four times, once for each calorimeter arch. Teams of student and post-docs from the collaboration were organized to assist in the final detector installation from August to November 2004.

Detector performance has been exceptional, more than 99.8 percent of the channels have been functional at all times since installation. The expected improvement in occupancy has been realized. The improved segmentation is being used for new detector algorithms that could not be attempted before.

### 1.2.2 Electromagnetic Timing

Collaborating Institutions:
- Texas A&M University
- University of Chicago
- INFN- Frascati
- Argonne National Laboratory

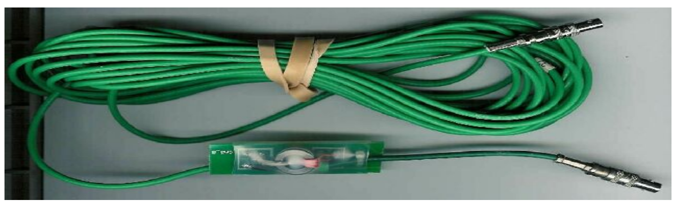

**Electromagnetic Timing Signal Splitter**

Timing readout for the electromagnetic calorimeters was recently installed as part of the upgrade to the Run II version of CDF. This system, known as EMTiming, is similar to the previous hadronic calorimeter system but has a resolution of less than a nanosecond and covers the central ($|\eta|<1.1$) and plug ($1.1<|\eta|<2.1$) portions of the calorimeter, where $\eta = -\ln(\tan(\theta/2))$, and $\theta$ is the angle from the beamline.

The design of the EMTiming system is optimized for searches for production of new particles that decay into high energy photons as would be the case in some models of Supersymmetry or Large Extra Dimensions. Final state particles from all proton anti-proton collisions typically traverse the CDF detector and the EMTiming system records the time of arrival for those that deposit large amounts of energy in the EM calorimeter. To improve the search sensitivity and robustness, the system can verify that all photons (or other energy) are from the primary collision and reject and estimate the rate of cosmic ray and beam-related backgrounds. In addition, the system also allows for a new class of searches for heavy, long-lived particles that decay to photons via their delayed time of arrival.

Particles from the collision that deposit energy in the EM calorimeters create light within the scintillators that can be used for a timing measurement. Photo-multiplier tubes (PMTs) collect this light and effectively convert the particle energy into an analog signal. Prior to the installation of the EMTiming system this signal was only used as part of the energy measurement. The EMTiming system routes a copy of the PMT signal to a passive Transition Board (TB) and Amplifier-Shaper-Discriminator board (ASD) pair that, in turn, sends their signal to a fixed-start Time-to-Digital Converter (TDC) board for a timing measurement.

In order to continue to use the pre-existing CEM hardware, an inductive signal ”splitter” board was designed that is placed between the PMT base and the original 25-foot RG174 cable that carries the anode signal for an energy measurement. The splitter effectively routes a fractional portion of the PMT pulse energy for timing use, while not affecting the energy measurement. The splitter solution avoids potential ground-loop problems since there is no electrical contact between the timing and energy readout lines.

The secondary pulse used for the timing measurement has a voltage that is 15 percent of the primary signal. Since the two lines are only inductively coupled, and the energy measurement is done with a charge integration device, in principle this solution should not affect the energy measurement since no charge can be lost. To test this, test stand comparisons of the integrated charge for a PMT pulse, with and without a splitter, were performed at various points over the full energy integration range. There was no observed (systematic or otherwise) effect on the linearity or resolution for all energies, with a measurement uncertainty of approximately 10 percent of the one sigma variation in the charge integration measurement itself, for a given energy.

The EMTiming system provides timing readout for the CDF electromagnetic calorimeters with good uniformity and resolution.[2] It has its 50 percent efficiency points at $3.8 \pm 0.3$ GeV and $1.9 \pm 0.1$ GeV in the CEM and PEM respectively, and is 100 percent efficient well above threshold. After a full set of corrections, $600 \pm 10$ ps and $610 \pm 10$ ps timing resolutions are found, respectively, with only small deviations from a Gaussian distribution.
There are very few pathologies observed in the data such as recording a time when no energy is deposited or recording multiple times for a single energy deposit. The system is well understood and ready for searches for new particles that decay to photons.

### 1.3 Data Acquisition and Trigger Upgrades

Several upgrades to the trigger and data acquisition systems were made as part of the Run IIb CDF Detector Project. These upgrades were all motivated by a need to maintain viability of the experiment in an environment of high luminosity. Two effects drove these subprojects. First, the rate and volume of data that must be processed and recorded posed a limitation on CDF’s effectiveness at high luminosity. Early projections indicated that a rate limitation of approximately 300 Hz existed in the original system for the rate of Level 2 triggers. A need for 1 kHz was projected in this parameter to maintain the physics program of the experiment. Secondly, high luminosity operation produces an event environment that contains multiple simultaneous collisions. Higher charged particle multiplicities that result will increase data volume, and add complexity to the events that must be processed by the track trigger systems. These trigger systems will tend to select an increasing number of fake triggers as the multiplicity rises. Upgrades to the track triggers were built to reduce the fake rates and processing times of these systems.

### 1.3.1 Tracking Chamber Time to Digital Converter Upgrade

Collaborating Institutions:
- University of Chicago
- Fermilab
- Duke University
- University of Michigan
- Texas A&M University

The Time to Digital Converters (TDCs) that are used to readout the Central Outer Tracker (COT) at CDF were identified as a system with an inherent rate limitation of about 300 Hz. The signal processing on the board, and the communications with the data acquisition system presented problems for high luminosity operation. Consequently, a new TDC was designed to overcome these limitations.

The new device is shown in the photograph below. The TDC was based on programmable gate arrays, so most of its functionality was programmable. Also, the design holds many fewer components than the original TDCs. Trigger information is provided on a backplane connector that is compatible with the existing system. The full performance characteristics and a detailed discussion of the Run IIb TDC has been published elsewhere[3].

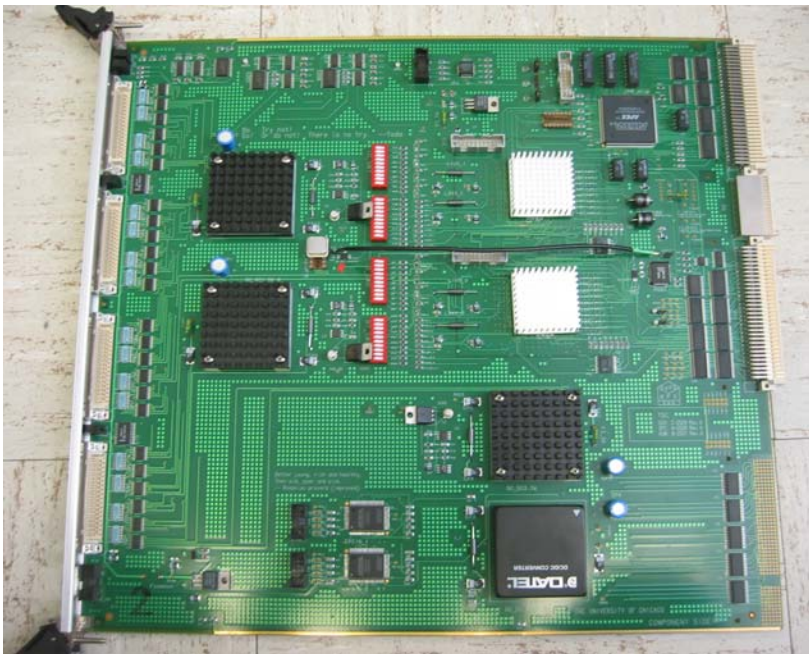

**The Run IIb TDC for the COT**

At the time of the Production Readiness Review for the Run IIb TDC, excellent performance of the new device had been demonstrated. Readout speeds of 2 kHz looked to be achievable, which was well in excess of the specification. Interesting developments had also taken place in the understanding of the older devices as well. The work on the new TDC motivated improvements to the data format used by the older device, which allowed for reduced data volume and increased throughput. Also, programming changes were made which decreased the time required for processing on the TDC itself. Finally, a relatively minor hardware change was identified, that would allow operation of the older devices at Run IIb rates. The review concluded that the modifications made to the older devices would yield a system of TDCs that would meet the Run IIb operations specifications. Project management concluded that this latter option would be pursued, to decrease schedule risk due to commissioning a new system. This change also resulted in a reduced cost to the project. The older TDCs were modified and reprogrammed in 2004-05 and are currently included in operations with the full set of Run IIb changes in place.

### 1.3.2 Level 2 Trigger Infrastructure Upgrade

Collaborating Institutions:
- Fermilab
- University of Chicago
- University of Pennsylvania

The CDF trigger is a three level system. The Level 1 and Level 2 triggers use custom designed hardware to find physics objects in a subset of the event information, while the Level 3 trigger uses the full set of information for complete event reconstruction in a computer farm. The Level 1 system has a synchronous 40 stage pipeline. When an event is accepted by the Level 1 trigger, all data are moved to one of the four Level 2 data buffers in the front end electronics, while a predefined subset of these data is sent to the asynchronous Level 2 trigger where a limited event reconstruction is performed and a Level 2 decision is evaluated. The Level 2 latency, the time to load and process the data for each event at Level 2, is critical to the throughput of the system. Deadtime arises when an event is accepted by the Level 1 trigger while all four Level 2 data buffer are occupied. To minimize the deadtime, both the loading and processing time has to be minimized. Loading time refers to how long it takes to deliver data to the Level 2 decision CPU after a Level 1 accept. Processing time refers to the time it takes for the CPU to unpack the data, to form objects, and to make a Level 2 decision.

The original Level 2 trigger was designed and built in the mid to late 1990's based on the technology available at that time. The design relied on a custom bus (Magic bus), a now obsolete processor (DEC Alpha) on a custom motherboard, and a set of different interface boards. The loading time and processing time are large with long tails, and the system was able to handle a Level 1 accept rate up to 25 kHZ (typically ran at luminosity below $1.2 \times 10^{32}$ cm$^{-2}$ sec$^{-1}$) with a Level 2 accept rate around 300 Hz. At the expected peak RunIIb luminosity of up to $3 \times 10^{32}$ cm$^{-2}$ sec$^{-1}$, close to 10 interactions per crossing will be observed. This implies that the average data size will increase substantially, and as a result, both the loading and processing time at Level 2 will grow significantly. This was the main motivation of the Level 2 decision crate upgrade. In addition, the concern for the long term maintenance of the existing system was another motivation for the upgrade. The easiest way of providing long term maintenance is to minimize the number of different types of custom boards and to rely on commercial or externally-supported components in as many places as possible.

The strategy for the new upgrade system is to convert and pre-process all Level 2 trigger data fragments from upstream to a self-describing format using a universal interface board, the so called "Pulsar" board, which is designed to be a general purpose interface VME board for HEP. Multiple copies of the board running different types of firmware are employed to process the data from all of the upstream subsystems. The data fragments are then merged (using the same type of Pulsar board) and transfered via a high speed common data link—S-LINK, designed for LHC experiments–into a modern Linux PC for decision making. The data transfer over PCI bus is done by using the high speed and high bandwidth S-LINK to PCI cards developed at CERN for Atlas experiment (the so called FILAR card). The CPU decision is then transfered over PCI bus (using the PCI to SLINK card from CERN) back to a Pulsar board which interfaces with the Trigger Supervisor.

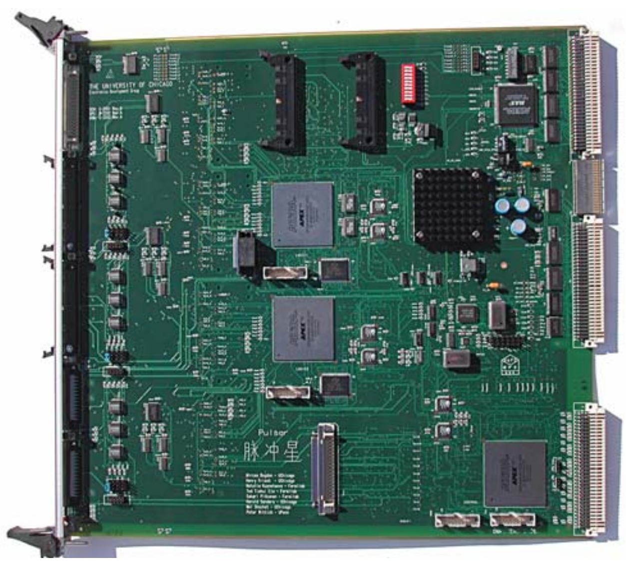


**The Run IIb Pulsar Board**

The core of this upgrade is the design and development of the Pulsar board. The design philosophy of the Pulsar was to use one kind of general purpose motherboard, with powerful FPGAs and SRAMs, to interface any custom data link with an industry standard link through the use of mezzanine cards. The motherboard was designed in such a way that it could be configured either as a data sink or data source. The board thus designed is fully self-testable, at board level as well as at system level. This important design feature makes it suitable to develop and tune an upgrade system in stand-alone mode, therefore reduce the impact on running experiment during the commissioning phase. A significant fraction of the Pulsar design effort was dedicated to extensive verifications by using the state-of-the-art computer aided design tools, which helped to streamline the design process significantly. This approach allowed us to build flawless prototype boards without the need of any revision, including a few different types of mezzanine cards and transition module.

The upgrade system configuration employs nine Pulsar boards, many of which use different mezzanine cards and FPGA firmware design. In addition to the decision node, the system also includes another PC which handles the task of communicating between the decision node and the rest of the CDF DAQ system. The universality of the Pulsar design allows us to test each data path, hardware as well as firmware, in a test stand using additional Pulsars configured in transmitter mode. All hardware and firmware were tested extensively in this controlled environment before being integrated into the system. For all Pulsar boards used in the system, diagnostic DAQ buffers have been implemented in an uniform way, allowing us to readout the intermediate information, data as well as timing, into the data stream. This design feature is essential for commissioning, optimizing, and long term maintenance of the system. In order to minimize the impact of the operation of the CDF experiment during the system commissioning phase of the Level 2 upgrade, all input data paths were split so that a copy of the input data was made available to the new system and allowed commissioning of the new system in parasitic mode this way. The new system was tested extensively using this methodology before dedicated beam time was requested to allow the new Level 2 system to drive the DAQ system. In fact, the Pulsar-based Level 2 trigger system worked on the first attempt in the test run with dedicated beam time.

The new system has been officially used for Level 2 trigger decision since March 2005. The subproject was successfully finished on schedule and on budget. The overall performance of the upgraded Level 2 decision crate has been shown to be much better than the old system. The new system has been very reliable, without one single Pulsar board hardware failure since installation. The Level 2 latency, defined as the time between the Level 1 Accept and the broadcast of the Level 2 decision, has been reduced by almost 20 microseconds on average. Together with the Level 2 SVT upgrade, the system now can run with Level 1 rate up to 35 kHZ with deadtime about 10 percent from Level 2 latency. This represents a significant increase in the physics capability of the CDF experiment, especially at higher luminosities in the future.

In summary, the CDF Level 2 decision crate was upgraded. The design of the upgrade system departs significantly from the previous implementation. It makes use of Pulsar board, a general purpose 9U VME interface board developed for HEP application, and CERN S-LINK technology, as well as an easily upgradeable commodity CPU to run decision algorithms. The new system is designed to have a safety margin both in performance and flexibility to meet the RunIIb trigger challenges and to use built-in test capabilities to speed up the commissioning process and to ease the long term maintenance effort. The Pulsar design is so general that it has also been used extensively in the SVT and XFT upgrade at CDF, where more than 40 Pulsars, with different firmware and mezzanine cards, are used. In fact, Pulsar boards are now also used outside CDF and HEP, for TOTEM experiment at LHC and for MAGIC experiment for astrophysics. The Level 2 decision crate upgrade is a project where the S-LINK technology developed at CERN for the LHC experiments is used for the first time in a high rate hadron collider environment.  Knowledge gained by using S-LINK at CDF is transferable to and back from the LHC community.

### 1.3.3/1.3.11  Track Trigger Upgrade

Collaborating Institutions:

- Ohio State University
- University of Illinois
- INFN – Pisa
- Purdue University
- Baylor University
- University of California, Davis
- Fermilab

Excellent central tracking is one of the cornerstones of the CDF experiment.  In the trigger, the identification and precise measurement of track parameters is performed by the Extremely Fast Tracker (XFT).  Tracks identified by the XFT are utilized in identifying the signatures of more than 75 percent of the data written to tape.  Projections based on early Run II data showed that operation at high luminosity (luminosity beyond the original Run 2a specification) would severely reduce the ability of the XFT to reject fake track patterns.  Consequently, a plan to upgrade the system, and provide greater fake rejection, was implemented.  The new strategy involves inclusion of more information for each track candidate, in the form of additional COT information.  The stereo layers of the COT were not included in the original XFT design.  The Run IIb upgrade project supplements the existing XFT system with higher precision tracking information from the outer three stereo layers.  The existing XFT system, based entirely on axial layers, is retained.  By choosing a design which supplements the existing system, the upgrade was able to be upgraded without interfering with operations of the existing system.

Using Run II data, a detailed emulation of the upgrade electronics was developed.  This emulation shows that fake track rates will be reduced by as much as 80 percent at peak luminosity, allowing the CDF experiment to acquire more data with higher purity.  The upgrade is beneficial even at lower luminosity, as it allows trigger bandwidth to be utilized on higher purity signal events.

Several hardware components were constructed for this upgrade.  New daughter boards (XTC II) were built to extract the trigger information from the TDCs.  These were commissioned and installed during the TDC modification procedure.  Data provided by the XTC II boards are transferred through a new transition card located on the rear of the crates that hold the TDCs.  Data is transferred from the detector to the CDF trigger rooms through a new set of optical fibers.  The input information is processed in “Finder” boards, which locate track segments in the stereo superlayers.  The Finder output is then transferred by fiber optic interface to Stereo Linker Association Modules (SLAMs) where the stereo segment information is merged with the axial tracking information from the existing XFT system.  Data from the Finders are also sent via fiber optic directly to the CDF Level 2 Pulsar trigger system, where additional processing power exists to further utilize stereo tracking information to more precisely identify trigger objects such as electrons and muons.

Seven flavors of custom boards were designed and constructed for this system. The parallel nature of this system permits large data throughput with excellent processing efficiency. The integrated system is capable of processing approximately 30 Gbytes/s of COT data. Flexibility is maintained through the extensive use of field programmable gate array (FPGA) technology. Common fiber optic input/output formats are used throughout the system.

The hardware is fully installed for this system and all of the individual components have been thoroughly tested. The integrated system is being commissioned and will be included in operations when the Tevatron resumes the 2006 run.

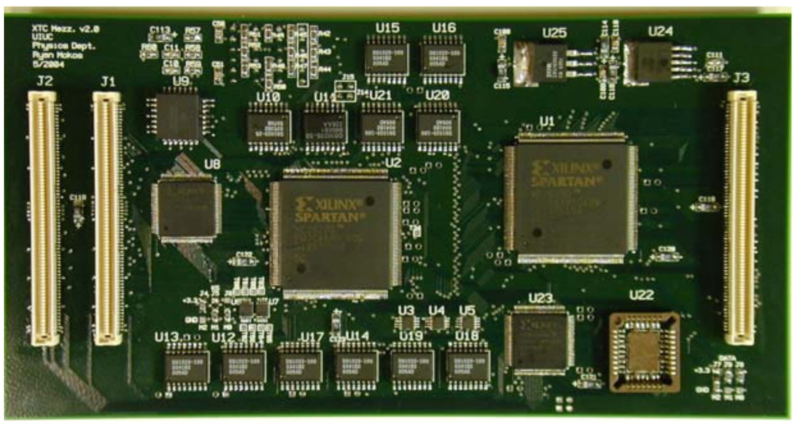


**The XTCII card**

**One crate of Stereo Finder Boards**

### 1.3.4 Event Builder Upgrade

Collaborating Institutions:
Massachusetts Institute of Technology
Fermilab

Part of the CDF Run 2b upgrade to cope with higher luminosity was an upgrade to the Event Builder (EVB) system, which is responsible for combining data from separate, independent front end systems into a full event for further trigger processing by the Level 3 system. The Event Builder system in use at the beginning of Run II operations had a rate limit of approximately 300 Hz, insufficient for the needs of the experiment at high luminosity. The system was replaced with an upgraded Event Builder, that is based on common network and computing technology. The upgrade system makes use of fast Ethernet for data transfer, and fairly standard single board computers for data acquisition.

A number of proprietary systems were eliminated, which leads to an overall reduction in maintenance cost over the duration of operations. In addition to using technologies that are fairly standard, this new system was built for a modest cost, compared to its predecessor. The upgraded event builder was installed in the summer of 2005, and was included into operations by September of that year. Routine operations in early 2006 saw event rates in excess of 600 Hz. The full specification will be achieved when operations resume later in the year.

### 1.3.5 Level 3 Trigger Upgrade

Collaborating Institutions:
- Massachusetts Institute of Technology
- Duke University
- Fermilab

The final trigger decision for CDF data is made in the Level 3 system. This trigger is a software system, which runs a subset of the offline algorithms, and selects events to be recorded. Increasing event rate and complexity increases the demands on the computing resources needed for this system. The Run IIb project acquired 320 dual processors PCs for the trigger decision and 32 additional personal computers to serve the input/output needs of Level 3. These machines are a standard architecture and packaging, which allowed for the most favorable price/performance. The newly acquired machines will be included in 2006 operations.

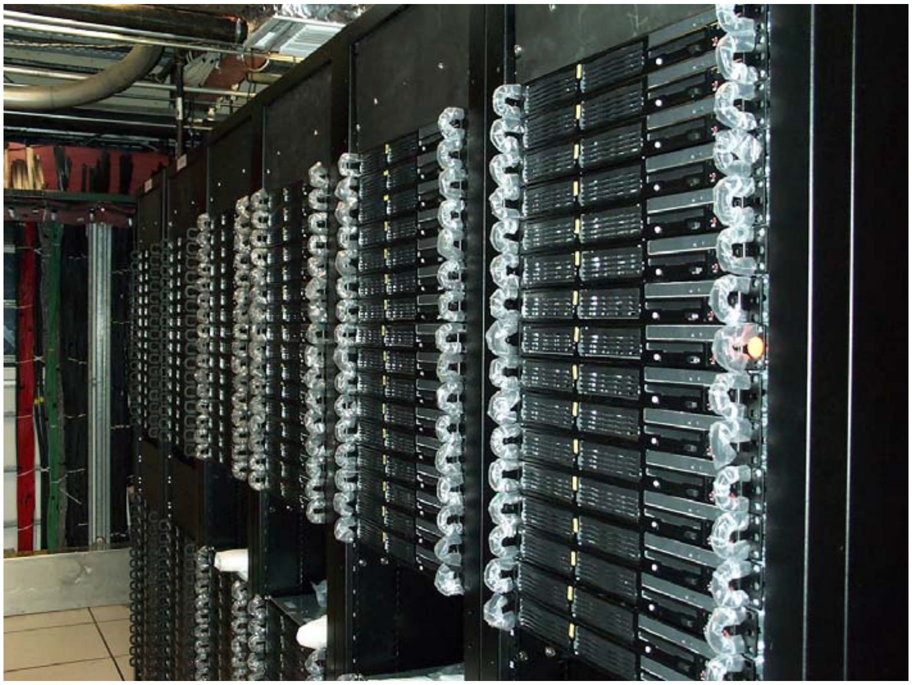

**Run IIb Level 3 processors**

### 1.3.6 Silicon Vertex Trigger Upgrade

Collaborating Institutions

INFN – Pisa
University of Chicago
Fermilab

The Silicon Vertex Trigger (SVT) is part of the Level 2 trigger of CDF II. The SVT reconstructs tracks by associating Silicon detector hits to Central Outer Tracker tracks reconstructed by the Level 1 trigger. The SVT improves tracking quality with respect to the Level 1 trigger. In particular, the SVT measures the transverse impact parameter with a resolution sufficient to allow selection of secondary vertices from B meson decay. In order to do this, the SVT must perform track fitting using the full resolution of the Silicon detector. The main challenge of the SVT is to perform track reconstruction within about 20 to 40 microseconds as required by the specifications.

The SVT required an upgrade in order to perform track reconstruction within the allowed 20 to 40 microseconds at the highest expected Tevatron luminosity, $3 \times 10^{32}$ $cm^{-2}$ $sec^{-1}$. In fact the SVT processing time is approximately proportional to the number of track candidates to be fit. As the luminosity increases, the occupancy in the CDF detector increases, and as a result the number of track candidates also increases. This results in a longer SVT processing time that would cause trigger dead-time or a reduction in the available Level 1 trigger bandwidth, thus highly reducing the CDF physics output.

In order to speed up the execution of SVT two actions were taken:

- Improvement of the resolution at the pattern recognition level. With higher resolution, fewer track candidates are found and fewer fits have to be performed.
- increase of the clock speed of the Track-Fitter board and of other interface boards, allowing faster track fitting and faster IO speed.

The pattern recognition is performed within a custom Associative Memory (AM). In order to improve the operational resolution of the AM without losing too much efficiency, a new AM system had to be designed and built. The new system allows 512,000 tracking patterns for each silicon detector wedge instead of 32,000 patterns. The larger number of patterns allows for thinner track patterns, thus improving the system resolution and reducing the number of fake track candidates. This required the design of a new AM board, called AM++, and of a new AM chip. The most important decision was to use standard cell technology for the new AM chip rather than full custom chips or FPGAs. The standard cell technology was the best compromise to achieve a high logic density (4000 patterns/chip) with a short design time of about one year. The easier FPGA option would not reach the same logic density. This part of the project was funded by INFN.

In order to increase the clock speed of other boards in the system, Pulsar boards were used. The choice of the Pulsar board significantly reduced the development time to about one year.  Since the SVT function requires much more memory (RAM) than available on the Pulsar, two custom memory mezzanines were designed to extend Pulsar memory.

For the AM upgrade 30 AM++ boards were built.  Each board includes 64 AM chips. Thirty Pulsar boards were bought to implement the TrackFitter (TF++) and HitBuffer (HB++) functions.  The AM-Sequencer-Road-Warrior (AMSRW) function was implemented on previously purchased Pulsar boards.  There were 64 RAM mezzanines of 4Mx48 bits and 64 RAM mezzanines of 512kx24 bits produced.

Hardware production and firmware development were the two major tasks.  In addition a lot of work went into integration with the existing system.  This consisted of software development for the initialization of the boards, parasitic testing of the boards, and development of monitoring and simulation software.  How to take advantage of the higher number of patterns available was studied.  In the first and current implementation they are being used to improve pattern recognition resolution and hence speed, but other applications are possible.

The system was installed in two steps as the boards became available.  During summer 2005, 12 AMSRW boards were installed, 12 AM++ boards (half loaded) and 12 TF++ boards.  During February 2005, 12 HB++ were installed and the AM++ system was completed to a total of 24 AM++ (fully loaded).  The system is now fully operational.

The attached plot shows the impact of the SVT and level-2 upgrades on the available Level 1 bandwidth.  The bandwidth is increased from 18 kHz to 25 kHz even when the instantaneous luminosity is doubled.  It is still 25 kHz at current peak luminosity thanks to the fact that the SVT processing time now has only little dependence on the event complexity.  The bottom line is that SVT triggers are kept operational at the highest luminosity.  At lower luminosity, more Level 1 bandwidth is obtained, i.e. more valuable data to tape.

## Cost and Funding

This project was completed under budget and ahead of schedule.  The final DOE MIE cost for the project was $7.2 million.  The project was able to achieve its mission need and also return $1.0 million of unused contingency, in addition to $2.2 million that was returned in June 2005 to the Office of High Energy Physics for other uses.  Essentially all costs were paid by May 2006, well ahead of the CD-4 date of November 2006.

| Run IIb CDF Detector Project Cost Table | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 2003 Baseline | | 2004 Baseline | | 2005 Baseline | | Final Cost | |
| | Total | DOE MIE | Total | DOE MIE | Total | DOE MIE | Total | DOE MIE |
| Silicon | 15.2 | 11.5 | 4.7 | 2.5 | 3.5 | 1.3 | 3.5 | 1.3 |
| Calorimeter | 1.1 | 0.3 | 1.0 | 0.3 | 1.2 | 0.5 | 1.2 | 0.5 |
| DAQ/Trigger | 4.7 | 4.0 | 4.1 | 3.8 | 5.6 | 4.7 | 5.5 | 4.7 |
| Administration | 1.6 | 1.3 | 1.0 | 1.0 | 0.7 | 0.7 | 0.7 | 0.7 |
| Contingency | 7.9 | 7.9 | 2.7 | 2.7 | 0.9 | 0.9 | 0 | 0 |
| Total Cost | 30.4 | 25.0 | 13.5 | 10.4 | 11.9 | 8.2 | 10.9 | 7.2 |

| Run IIb CDF Detector Project Funding Table | | | | |
|---|---|---|---|---|
| | 2003 Baseline | 2004 Baseline | 2005 Baseline | Final Cost |
| DOE MIE | 25.0 | 10.4 | 8.2 | 8.2 |
| DOE R&D | 2.1 | 2.1 | 2.1 | 2.1 |
| Foreign | 2.9 | 0.6 | 1.2 | 1.2 |
| US Universities | 0.4 | 0.4 | 0.4 | 0.4 |
| Total | 30.4 | 13.5 | 11.9 | 11.9 |

**Major Baseline Change Control Log**

2003 Baseline: This baseline was the original baseline.

2004 Baseline: Due to circumstances outside of the project's control, the scope of the project was reduced. Accelerator performance projections were reduced in 2002 and 2003, which motivated the Fermilab Director to recommend cancellation of full production of upgraded silicon detector. Consequently, the project scope was reduced in December 2003 to eliminate the construction of the silicon detector except for the construction of approximately seven percent of the silicon detector elements to serve as a demonstration device.

2005 Baseline: In 2005 it was evident that the project would be completed significantly under budget. Therefore, the project funding was reduced by $2.2 million for the Office of High Energy Physics to fund other needs in fiscal year 2005.

**Schedule**

The Project's Level 0 Milestones are as follows:

| | Milestone Description | Baseline | Actual |
|---|---|---|---|
| CD-0 | Approval of Mission Need | May 2001 | May 2001 |
| CD-1 | Approve Preliminary Baseline | Feb 2003 | Feb 2003 |
| CD-2 | Approve Performance Baseline | Feb 2003 | Feb 2003 |
| CD-3a | Approve Limited Construction | Feb 2003 | Feb 2003 |
| CD-3b | Approve Full Construction | Dec 2003 | Dec 2003 |
| CD-4 | Completion of Construction | Nov 2006 | |

Below are lower Level 1 milestones that were established in the Project Execution Plan to monitor project performance. Each milestone was accomplished on schedule.

| Milestone | Baseline | Actual |
|---|---|---|
| Calorimeter Upgrades Ready for Installation | Jan 2006 | Jan 2005 |
| Data Acquisition and Trigger Upgrades Ready for Installation | Jan 2006 | Dec 2005 |

**Project Management**

Fermilab performed project management under the auspices of the DOE Chicago Operations Office. The Fermilab Project Manager was Patrick Lukens with the Collider Detector at Fermilab (CDF) Department of the Particle Physics Division (PPD). The DOE Project Director was Paul Philp of the Fermi Site Office (FSO). The Program Manager was Michael Procario of DOE's Office of High Energy Physics.

The organization chart of the project at the time of its Critical Decision 2 and 3a is shown below. Project Management Group meetings were conducted on a monthly basis to track progress, evaluate change control requests, and provide general guidance on projects activities. Scientific guidance of the project was provided from the collaboration, through the spokesmen of CDF. Over the course of the project, the individuals in several of the positions within the organization changed. However, the structure was held for the duration of the project.

## DOE Organization

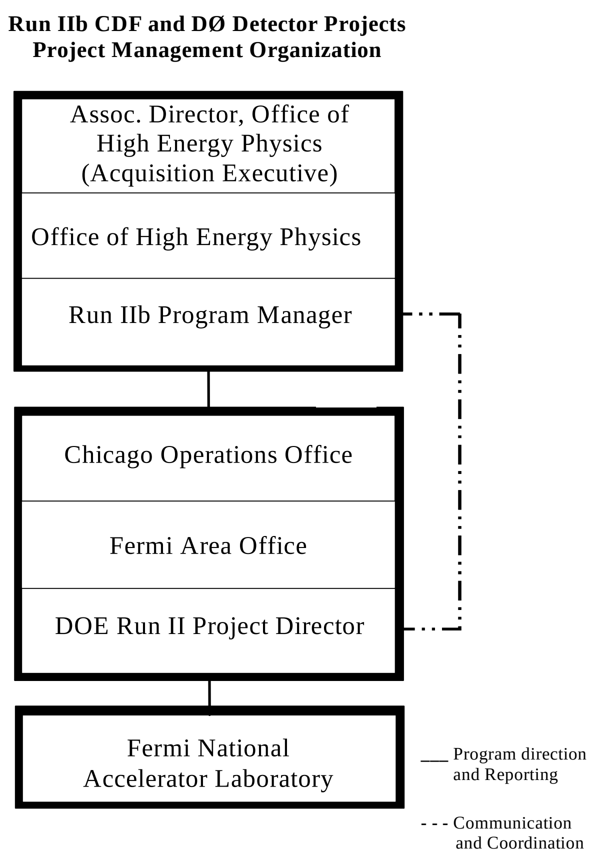

## CDF Project Organization (Fermilab)

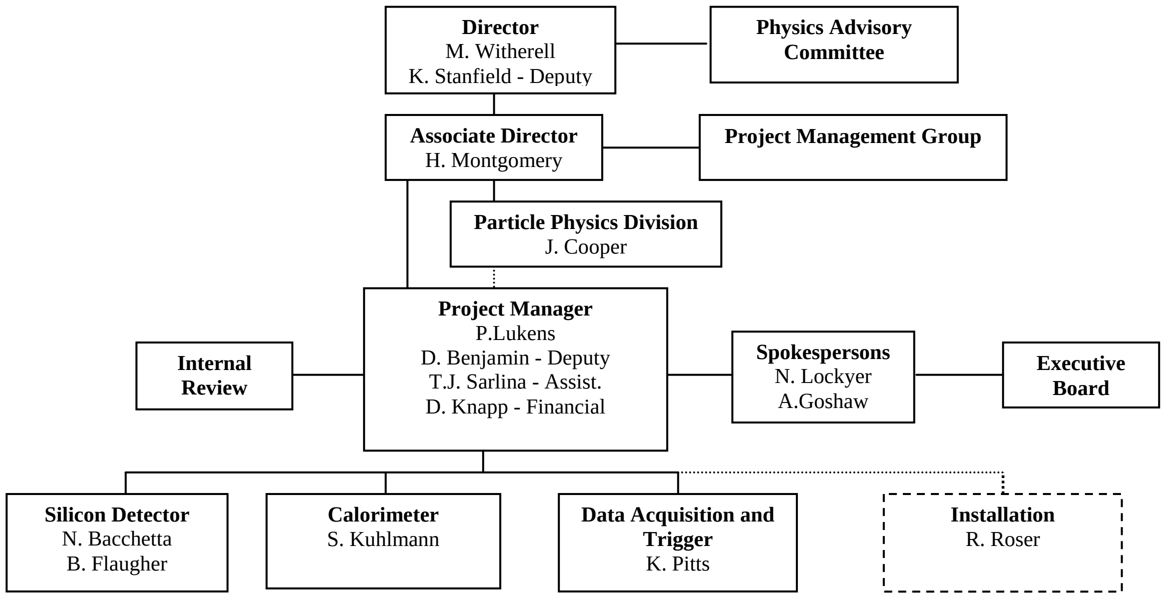


## Project Tracking

Several formal reviews of the project were held both before the establishment of CD-3, and during the project's execution. A list of the major reviews is given in the table below:

| Review | Date |
|---|---|
| Physics Advisory Committee | November 2001 |
| Director's Technical Review Committee | December 2001 |
| Physics Advisory Committee | April 2002 |
| Director's Review | April 2002 |
| Physics Advisory Committee | June 2002 |
| Director's Review | August 2002 |
| Office of Science Baseline Readiness Review | September 2002 |
| DOE External Independent Review | November 2002 |
| P5 Subcommittee of HEPAP | March 2003 |
| Director's Review to Establish a New Baseline | November 2003 |
| Director's Review | July 2004 |
| Director's Review | January 2005 |

Director's Reviews were organized by the Project Management Office within the Fermilab Directorate. In addition, meetings of the Project Management Group were held on a monthly basis, and the DOE Project Director and Fermilab Project Manager held weekly meetings. The Project issued monthly reports on both technical, financial, and schedule status to the laboratory management and DOE Project Director. Earned Value was reported prior to the baseline change in 2003. However, the total cost of the project did not warrant this reporting after cancellation of the silicon detector production, and the reporting was discontinued. The DOE Project Director produced quarterly reports, participated in quarterly status meetings with SC, and provided monthly status updates to the Project Assessment Reporting System.

Internal reviews of specific subprojects were held as needed. Most subprojects held a Production Readiness Review when they were ready to place procurements for production quantity parts. These reviews were designed to assure that sufficient quality control was in place, before committing significant resources. All subprojects were also given Installation Readiness Reviews, which were coordinated by the Head of Operations at CDF. These reviews assured that the installation and commissioning strategy adopted would produce a successful result consistent with operations.

## Safety

The equipment constructed by the project does not have an impact on the safety of CDF. No additional hazards were introduced by the project. Consequently, no changes to the CDF Safety Assessment Document were made or required. Although not strictly within the scope of the project, installation of the various components had the most significant safety implications, because of the work environment in the CDF collision hall, and the protections needed to work safely. The most challenging installation was the Central Preshower Detector. This installation was reviewed by CDF operations and Fermilab ES&H staff before the work was allowed to begin. All personnel were trained in the hazards, and the work proceeded successfully, and without any accidents. Other installations are more routine, but proper training is always required for these installations and collision hall work, as well.

## Lessons Learned

The lessons learned from the project include:

- The project benefited from being flexible to changes in operational needs. Due to changes in the projected deliverable luminosity, the silicon detector portion of the project was deemed to be unnecessary. This was a significant change to the project. The project participants rallied and accomplished the project missions need.
- The project benefited from including Value Engineering throughout the life of project. Several subprojects, e.g. the TDC subproject, benefited by value engineering during the execution phase of the project. This resulted in a reduction in cost for some subprojects, a reduction in the required variety of parts that were needed, and an improvement in installation strategies.

---

[1] T. Akimoto *et a*l., Nucl. Inst. and Meth., A556 459-481 (2006)
[2] FERMILAB-PUB-05-54 3-E, Dec 2005
[3] Bogdan *et. al*, Nucl. Instr. and Meth., **A554**:444-457, 2005

Project Completion Sign-off Sheet

The Run IIb CDF Detector Project has been successfully completed and has met the requirements as stated in the Project Execution Plan for CD-4 milestone: Approve Project Completion.

Patrick T. Lukens — date
Project Manager

Jaco Konigsberg — date
Co-Spokesperson

Robert M. Roser — date
Co-Spokesperson

Hugh E. Montgomery — date
Associate Laboratory Director

Paul R. Philp — date
DOE Project Director

Joanna M. Livengood — date
Fermi Site Office Manager

Michael P. Procario — date
Program Manager